arXiv:1601.07613v1 [cs.DC] 28 Jan 2016

# Edge coloring in unstructured CFD codes

Andrew Giuliani Lilia Krivodonova

**Abstract**

We propose a way of preventing race conditions in the evaluation of the surface integral contribution in discontinuous Galerkin and finite volume flow solvers by coloring the edges (or faces) of the computational mesh. In this work we use a partitioning algorithm that separates the edges of triangular elements into three groups and the faces of quadrangular and tetrahedral elements into four groups; we then extend this partitioning to adaptively refined, nonconforming meshes. We use the ascribed coloring to reduce code memory requirements and optimize accessing the elemental data in memory. This process reduces memory access latencies and speeds up computations on graphics processing units.

## 1 Introduction

Graphics processing units (GPUs) are massively parallel platforms that have become useful in computational fluid dynamics (CFD) solvers. On such architectures, data are stored in shared memory and manipulated by processes that solve the problem in parallel. If multiple processes, or threads, write simultaneously to the same memory location, a race condition is created. Computations can then have unpredictable results in the sense that they become dependent on the order by which contentious memory locations are accessed [1]. Therefore, care must be taken to develop algorithms that do not lead to race conditions.

A race condition can arise in the evaluation of the surface integral in discontinuous Galerkin (DG) and finite volume type numerical methods. We will illustrate this issue on an example of the DG method used to discretize the conservation law

$$\frac{d}{dt}\mathbf{u} + \nabla \cdot \mathbf{F}(\mathbf{u}) = 0, \tag{1}$$

with the solution $\mathbf{u}(\mathbf{x},t) = (u_1, u_2, ..., u_M)^\intercal$, $(\mathbf{x},t) \in \Omega \times [0,T]$, and the flux function $\mathbf{F}(\mathbf{u})$. We divide the domain $\Omega$ into a mesh of elements, e.g., triangles, quadrilaterals, tetrahedra for two and three-dimensional problems. Discretizing (1) on this mesh with the DG method [2] yields the scheme

$$\frac{d}{dt}\mathbf{c}_{i,j} = \int_{\Omega_i} \mathbf{F}(\mathbf{U}_i) \cdot \nabla v_{i,j} d\Omega_i - \sum_q \int_{I_{i,q}} v_{i,j} \mathbf{F}(\mathbf{U}_i, \mathbf{U}_{p_q}) \cdot \mathbf{n}_{i,q} dI_{i,q}, \tag{2}$$

where the numerical solution $\mathbf{U}_i$ on the physical element $\Omega_i$ is approximated by a linear combination of $N_p$ orthonormal basis functions $v_{i,j}$, i.e. $\mathbf{U}_i = \sum_{j=1}^{N_p} \mathbf{c}_{i,j} v_{i,j}$ with

$\mathbf{c}_{i,j} = [c_{i,j}^1, c_{i,j}^2, \ldots, c_{i,j}^M]^\intercal$ as the modal degrees of freedom (DOFs). The numerical flux $\mathbf{F}(\mathbf{U}_i, \mathbf{U}_{p_q})$ is computed on the surface shared by adjacent cells $\Omega_i$ and $\Omega_{p_q}$, and $\mathbf{n}_{i,q}$ is the outward facing normal on the $q$th surface of element $i$. We refer to the elements that share the surface $e_k$ as the left and right elements of that surface. Depending on the dimension of the problem, surfaces of the element $\Omega_i$ can geometrically be either edges (of, e.g., a triangle) or faces (of, e.g., a tetrahedron). For simplicity of illustration, we will mostly discuss edges although the ideas are applicable to faces as well.

Integrating the numerical solution in time requires evaluation of the right-hand side of (2), which is composed of a volume integral over $\Omega_i$ and surface integrals over $I_{i,q}$. The volume contribution is easy to parallelize as it only requires information local to $\Omega_i$. However, the surface contribution requires more care as it involves writing data to memory locations for both $\Omega_i$ and its neighbors $\Omega_{p_k}$. The surface contributions can be computed in parallel by assigning one thread per edge in the computational mesh as shown in Figure 1. The processes $p_k$ and $p_j$ are then tasked with computing the surface terms along edges $e_k$ and $e_j$, respectively. Once this is done, the processes must save the result at memory addresses corresponding to the edge's left and right elements, e.g. $p_k$ writes the surface contribution to locations for $\Omega_0$, and $\Omega_1$, and $p_j$ writes to locations for $\Omega_1$, and $\Omega_2$. If both processes write simultaneously to the address for $\Omega_1$, a race condition will occur.

The mitigation of such a race condition can be done in a number of fashions. First, each process can be given a portion of buffer memory such that no two processes store their result at the same location. For example, process $j$ could write data for its left and right element to addresses $2j$ and $2j+1$, respectively. Subsequently, these data in the buffer can be combined using an additional parallel kernel [2, 3, 4]. The downside of such an approach is the extensive use of buffer memory for intermediate calculations. The length of these buffer vectors is the number of DOFs per element times the number of surfaces; we report the memory requirements of these buffers in Section 4.1. Atomic operations have also been proposed, but may degrade the solver's efficiency by an order of magnitude [2]. Finally, the element-wise surface integration approach has been discussed and shown to be suboptimal in [2, 5, 4], as fluxes at surface integration points are evaluated twice.

Another approach to avoiding a race condition is coloring the surfaces of the mesh. The surfaces of a mesh can be colored into separate groups such that no element possesses surfaces that belong to the same color. A sample two-dimensional mesh and a possible edge coloring are shown in Figure 2. Each edge has a color (1,2, or 3) and no triangle has multiple edges of the same color. Clearly, three is the minimum number of colors for this element geometry. Once the set of edges is partitioned into three groups, the surface contributions are evaluated color-by-color and added to the variable $\mathbf{c}^{\text{rhs}}$, which stores the volume contribution as shown in Algorithm 1. Because no edges in a color group share the same element, there will be no memory in contention, that is, the race condition will be avoided.

Edge coloring has been used in the context of MPI computing, for example [6], to distribute edges between cores. It has also been used in GPU computing, e.g. [7, 8], where they were not concerned with using the optimal number of colors. Typically, a naive greedy coloring algorithm will loop through all surfaces in the mesh. For each surface,

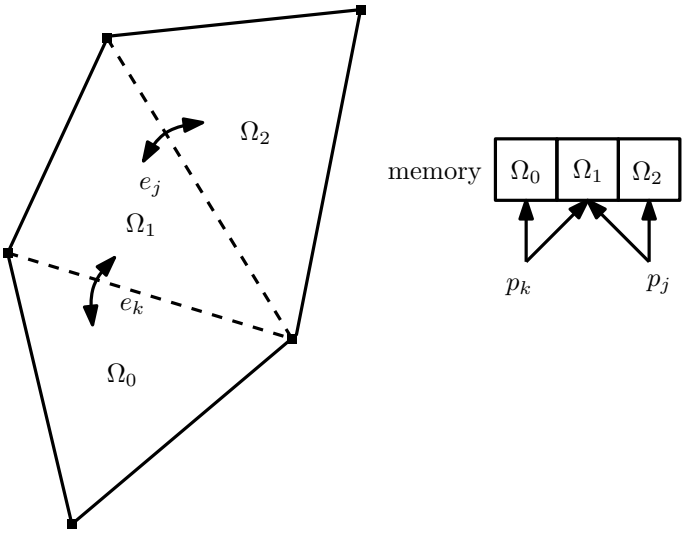


Figure 1: Sample mesh (left); processes $p_k$ and $p_j$ both write surface contributions to the memory location for $\Omega_1$ (right).

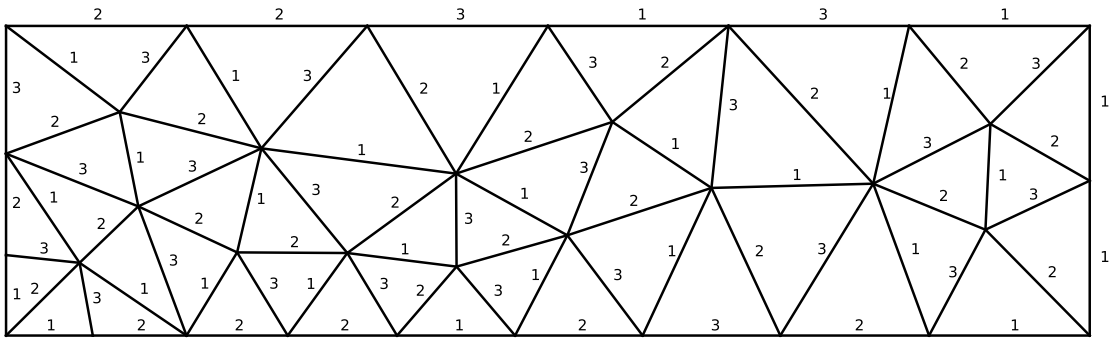

Figure 2: An example of edge coloring. Each triangle has edges of distinct, non-repeated colors 1, 2 and 3.

the edge colors of its the left and right element are checked. Then, the first available color is assigned to the surface. This greedy algorithm will yield meshes of a maximum five colors for triangular elements, and a maximum of seven colors for quadrilateral and tetrahedral elements. This is not the minimum number of colors possible for these element geometries. In this work, we propose an edge coloring heuristic that results in an optimal number of colors for a variety of element geometries. This allows for a reduced number of kernel launches and streamlining of the code. The overhead associated with launching extra kernels is minor for GPUs, but non-negligible for multi-core CPUs [7]. Our testing on GPUs has shown that using a suboptimal number of colors can decrease execution efficiency by approximately 5 %, which is negligible, see Table 7. However, an excessive number of colors can be a detriment to code simplicity. This is especially true with respect to adaptive mesh refinement, where the number of required colors can rapidly increase. For example, with tetrahedral elements a greedy coloring algorithm would yield seven colors versus our algorithm which yields the optimal four colors. Upon introducing refined elements whereby the difference in level of refinement between adjacent elements is one, these numbers increase to 16 and 28 colors, respectively. More aggressive refinement would increase the number of colors even more substantially. Therefore, for code simplicity and a reduced number of kernel launches, it is worthwhile to use the fewest possible number of colors to partition the set of edges in the mesh.

In order to improve the performance of unstructured CFD solvers, elements and sides may be ordered in memory to reduce memory access latencies, e.g., space-filling curves and the bin-ordering method have been used on single and multi-processor machines [6, 9, 4]. The speed-up observed depends on the hardware architecture [6] and software implementation. Various renumbering techniques are compared in [10] based on the performance of an edge-based solver for the Euler equations. A maximum speed-up of approximately 20 % was obtained compared to the ordered output from the mesh generator. In this work, we propose an ordering scheme based on edge coloring that exhibits comparable speed-ups for GPUs and show that using a minimal number of colors maximizes the attainable speed-up.

Thus, to eliminate memory contention, reduce memory requirements, and speed-up computations, we color the surfaces of the mesh into separate groups, e.g. colors 1, 2, and 3 for triangular elements and 1,2,3, and 4 for tetrahedra. The objective is to seek a coloring with the minimum number of colors. First, we apply a modified greedy coloring algorithm described in Section 3. This algorithm has linear complexity, but colors the mesh imperfectly, i.e., certain surfaces remain uncolored. We then post-process the created greedy coloring to ensure all surfaces are colored with the minimum possible number of colors, also described in Section 3. We numerically demonstrate that the coloring finishes in linear time on triangular, quadrilateral, and tetrahedral elements geometries.

# 2 Edge coloring

The coloring of the edges (or faces) of the computational mesh can be related to a standard edge coloring problem from graph theory for the element-wise connectivity

**Algorithm 1** Surface contribution evaluation

**Ensure:** the surface contribution stored in the variable $\mathbf{c^{rhs}}$.

```
procedure SURFACE EVALUATION
    for i ∈ {1, 2, ..., N_colors} do
        for all e_k in parallel do
            (l, r) ← the indices of two elements that share the edge e_k
            for j ∈ [1, ..., N_p] do
                c^rhs_{l,j} ← surface contribution for Ω_l along e_k
                c^rhs_{r,j} ← surface contribution for Ω_r along e_k
            end for
        end for
        synchronize threads
    end for
end procedure
```

graph. Given a computational mesh, we will now construct such a graph.

A computational mesh $(V, E, \Omega)$ can be described by its vertices $v_i \in V$, edges (or faces) $e_k \in E$, and elements $\Omega_i \in \Omega$. For a given mesh, we can construct an associated graph $G(N, L)$, by first placing a node $n_i \in N$ at the barycenter of the element $\Omega_i \in \Omega$. Further, two nodes $n_i, n_j$ are connected by a line $l_k \in L$ when the elements $\Omega_i$ and $\Omega_j$ are adjacent, i.e. they share an edge (or face). In other words, $G$ is a graph showing the connectivity between elements in the computational mesh. This process is illustrated in Figure 3.

The number of lines in $L$ incident on a node $n_i$ in $N$ is called the degree of $n_i$, $\deg(n_i)$. Vizing's theorem [11] states: a graph can be colored with $\max_{n_i \in N} \deg(n_i) + 1$ colors. That is, the connectivity graph of meshes with triangular elements can always be colored with four colors. Likewise, meshes with quadrilateral and tetrahedral elements can always be colored with five colors.

Furthermore, for planar meshes of triangles, it has been shown that only three colors are necessary [12]. The proof is based on assigning a color to each geometrical vertex $v_i \in V$ of the planar mesh of triangles. The colors must be placed such that no two vertices connected by a geometrical edge have the same color. By the four color theorem [13], this can be done with four colors if the mesh is a planar graph. From this vertex coloring, it can be shown that a valid edge coloring with three colors is always possible. This result also holds for meshes of triangles embedded on a sphere. Although it is not intuitive, such meshes are equivalent, i.e., isomorphic to planar graphs.

## 2.1 Approaches to coloring

The coloring algorithm resulting from the constructive proof of the four color theorem [13] is too complex for practical mesh coloring because it requires the handling of hundreds of reducible configurations. However, there exist heuristic algorithms that are simple and have been shown to terminate. The color exchange algorithm presented in [14] is a simple probabilistic algorithm that terminates in polynomial time for general graphs. It

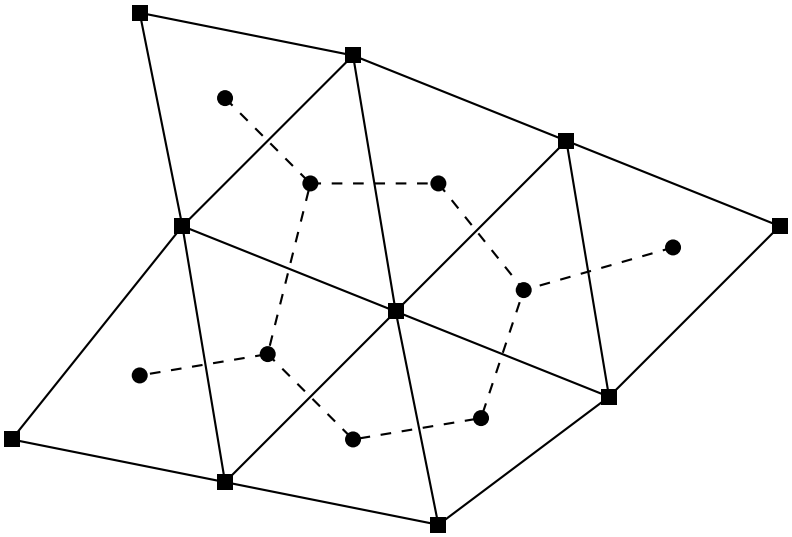

Figure 3: A computational mesh is shown with solid edges and square vertices. The mesh's associated connectivity graph has dashed lines and circular nodes.

is composed of two stages: first the mesh is colored imperfectly with a greedy algorithm, next a conflict resolution stage is completed to obtain an acceptable partitioning. During this last stage, the algorithm sometimes enters an unterminating loop. The loop can be broken by swapping colors in a prescribed manner. An alternative method is to restart the entire algorithm with a new initial greedy coloring in the hopes that a loop is not encountered again. This is the conflicting vertex displacement (CVD) algorithm for edge coloring graphs presented in [15]. We propose an algorithm with a simpler resolution of unterminating loops than the one presented in [14] that does not require recoloring the entire mesh like the one in [15].

# 3 Algorithm

Triangular and tetrahedral elements have 3 edges and 4 faces, respectively. Consequently, the smallest possible number of colors is three for triangles and four for tetrahedra. We call the set of available colors $C$, with $C = \{1, 2, 3\}$ for triangles and $C = \{1, 2, 3, 4\}$ for quadrangles and tetrahedra. We begin by describing the notation that will be used in the following subsections. The $k$th edge (or face) in the computational mesh is named $e_k$, and the $i$th element in the computational mesh is named $\Omega_i$. The two elements that share $e_k$ are denoted by $\Omega_{k_l}$ and $\Omega_{k_r}$, i.e. the left and right element of edge $e_k$. A nonconflicting coloring of an element $\Omega_i$ is such that the edge colors of $\Omega_i$ do not repeat. A nonconflicting coloring of $e_k$ is such that both $\Omega_{k_l}$ and $\Omega_{k_r}$ have nonconflicting colorings.

## 3.1 Conforming meshes

We now describe our coloring algorithm for conforming meshes. It is composed of two stages: a modified greedy coloring procedure, followed by a conflict resolution step. To

simplify the illustration, we describe the algorithm for triangular meshes. However, analagous arguments can be made for meshes of other element geometries, such as quadrangles, tetrahedra, etc.

At the start of the coloring algorithm, the color of each edge is initialized to -1, meaning that all edges have not been assigned a color. Then, the modified greedy coloring procedure passes through all edges $e_k$ in the mesh. For each edge $e_k$, the edge colors of its the left and right elements ($\Omega_{k_l}$ and $\Omega_{k_r}$) are checked. If possible, a nonconflicting color in $C$ is randomly assigned to $e_k$. A standard greedy coloring algorithm augments the set of colors when none are available to create a nonconflicting coloring for $\Omega_{k_l}$ and $\Omega_{k_r}$. Instead of increasing the number of colors in $C$, our modified greedy coloring algorithm leaves edges uncolored, i.e. of color `-1`, when there are no colors in $C$ that yield a nonconflicting coloring. For example, in step 1 of Figure 4, edge $CB$ cannot be assigned a color without causing either triangle $ABC$ or $CBD$ to have a conflicting coloring. Uncolored edges are called conflicts.

We now describe the procedure that seeks to resolve the conflicts created by the modified greedy algorithm, with the aim to obtain a nonconflicting edge coloring for all edges. We consider the uncolored edges one at a time. For an uncolored edge $e_k$, we exchange its color, i.e., `-1`, with the color of an edge that belongs to $\Omega_{k_l}$ or $\Omega_{k_r}$ such that the number of conflicts does not increase. For example, in Figure 4: step 1, the conflict on $CB$ may be swapped with the color of either edge $CA$ or $CD$. This operation does not increase the total number of conflicts in the mesh. However, swapping colors of $BA$ and $BC$ will create a new conflict as $BCD$ will have two edges with the same color, 3.

A simple geometric consideration reveals that a swap that does not increase the conflict count is always possible. After the swap, the conflict can either be resolved by choosing a nonconflicting color or it can be propagated further. For example, the conflict can be moved along the sequence of edges $BC$-$CD$-$DF$-$EF$ as illustrated in Figure 4: steps 1, 2, 3, and 4. After that, the conflict on edge $EF$ can be resolved by assigning it to color 2 (Figure 4: step 4). Conflicts generally cancel out one another, or resolve once they reach a boundary edge. Once a conflict is resolved, we move on to the next uncolored edge, i.e. conflict, and try to resolve it.

Some conflicts will be resolved in a finite number of swaps, and some will create a loop and visit the same sequence of edges. The latter case is illustrated in Figure 5 steps 1, 2, and 3. A loop is detected when the conflict is directed back to an edge that has already been visited. In this case, the algorithm attempts to break the loop by choosing a swap that increases the number of conflicts by one. For a conflict on edge $e_k$ that is in a loop, we uncolor two edges of $\Omega_{k_l}$ and $\Omega_{k_r}$ that share the same color, i.e. assign -1, to them. In step 4 of Figure 5, the conflict in a loop on edge $AB$ is moved to both $BD$ and $BF$ (two edges that previously shared the same color). Thus, the number of conflicts in the mesh is increased by one. The algorithm then attempts to resolve the next conflicting edge in the mesh.

The greedy coloring runs linearly in the number of edges as each edge is only visited once. We do not have a proof that the conflict resolution stage always terminates, so we cannot determine its theoretical complexity. However, the runtime seems to scale linearly with the number of edges in the mesh according to the timings in Section 4.

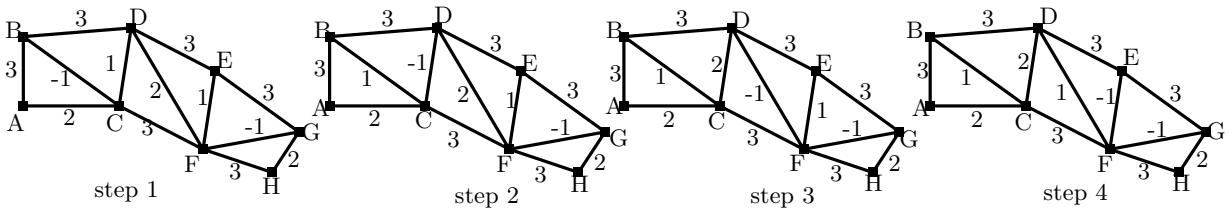


Figure 4: Both conflicts in step 1 can be resolved once the conflict originally on $BC$ reaches $EF$ in step 4.

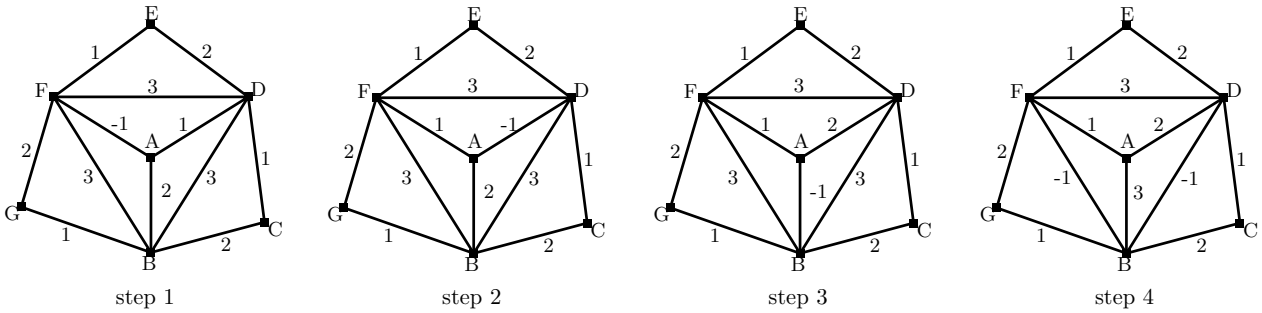


Figure 5: The conflict on $AF$ is stuck on a loop $AF$-$AD$-$AB$. Moving the conflict from that loop to $FB$ creates an additional conflict on $DB$ in step 4.

## 3.2 Nonconforming meshes

In this section we propose a coloring algorithm for nonconforming meshes, which are often a result of adaptive mesh refinement. We assume that computations start with a conforming mesh as most CFD codes do. The set of edges in the original mesh is partitioned using the algorithm described in Section 3. Based on this partition we color the refined mesh using the minimum possible number of colors. This process can be performed naturally in parallel on a GPU. The algorithm allows for easy transition between multiple levels of refinement without introducing conflicts. In this section, we make an assumption that adjacent elements differ by at most one level of refinement, though the idea can be extended to less smooth meshes.

### 3.2.1 Triangles

One way a triangle can be refined is by splitting it into four smaller triangles by connecting the midpoints of its edges as shown in Figure 6. This might produce a conconforming mesh if the adjacent element is not refined (Figure 6). We propose a coloring of these smaller triangles that doubles the number of colors initially present in the mesh to 6. This is the minimum possible number because a coarse element might have six refined neighbours (Figure 7). We note that some refinement strategies will refine elements with six neighbors to improve mesh smoothness, but for completeness we not do so here [16].

We call the edges of the original (or parent) triangle parent edges, each with a parent color $c^{(n)}$ with $n = 1, 2, 3$. The parent edges are ordered counterclockwise, e.g., the edges of triangle $DBA$ in Figure 6, left, are ordered $DB$, $BA$, and $AD$. Therefore, in this example the parent colors are $c^{(1)} = 2$, $c^{(2)} = 1$ and $c^{(3)} = 3$.

In Figure 6, right, triangle $DBA$ is refined. Each parent edge is divided into two child edges of equal length, e.g. the second edge $BA$ becomes $BG$ and $GA$. The first child inherits the color of the parent, and the second child's color is the parent's color shifted by three. We can write this formally with the mapping $c_m^{(n)} = [(c^{(n)} + 3m - 4) \bmod 6] + 1$ for the the $m$th child edge on the $n$th parent edge. The child edges are also ordered counterclockwise, e.g., the second parent edge $BA$ has child edges that are ordered $BG$, $GA$. Therefore, in this example the color of each child edge is $c_1^{(2)} = 1$, $c_2^{(2)} = 4$. Further, three additional edges are created in the element's interior: $GE$, $EF$ and $FG$. Each edge is prescribed the color of the parent edge to which it is parallel, e.g. $GF$ takes the color of $AD$ because they are parallel.

Coarsening the mesh consists of merging four small elements into one. In the coarsened triangle, a parent edge takes the color of the child edge to which it is parallel. That is, the coarsened triangle inherits the colors of the interior small triangle. This approach recovers the coloring of the parent edges before refinement.

### 3.2.2 Tetrahedra

A tetrahedron can be refined by splitting it into eight smaller tetrahedra, as shown in Figure 8. This is done by splitting each face into four triangles (Figure 9), and combining them with interior faces to form smaller tetrahedra. Coloring the child tetrahedra requires multiplying by four the number of colors initially present in the conforming mesh. 16 is the minimum possible number of colors because in the extreme case, a coarse tetrahedron can have 16 refined neighbors.

The four faces of the parent tetrahedron are called parent faces, each with a parent color $c^{(n)}$ for $n = 1, 2, 3, 4$. The parent faces of the tetrahedron $ABCD$ are ordered by our convention $ABD$, $BCD$, $CAD$, $CBA$.

The faces on the parent tetrahedron $ABCD$ in Figure 8 are each refined into four child faces. The color of the $m$th child face on the $n$th parent face is defined by the mapping $c_m^{(n)} = [(c^{(n)} + 4m - 5) \bmod 16] + 1$, where the child faces are ordered starting with the center face, then going counterclockwise from the bottom left, e.g., the refined faces of $ABD$ are ordered $HEI$, $AEH$, $EBI$ and $HID$. Therefore, the color of each child face is $c_1^{(1)} = 1$, $c_2^{(1)} = 5$, $c_3^{(1)} = 9$ and $c_4^{(1)} = 13$.

Further, eight faces are created in the interior of the parent tetrahedron. The first four, $GFJ$, $HEJ$, $IEF$, and $GHI$ (the grey faces in Figure 8 right) can take the color of the parent face to which they are parallel, i.e. $ABD$, $BCD$, $CAD$, and $CBA$, respectively. The final four interior faces are not parallel to any parent face, and are ordered $JIE$, $JIF$, $JIG$ and $JIH$. The color of the $n$th face in this list is defined by the mapping $c_n^{\text{interior}} = (c^{(n)} + 3) \bmod 16 + 1$.

As with triangular elements, it is straightforward to invert this mapping upon coarsening the eight child tetrahedra. The four child faces $GFJ$, $HEJ$, $IEF$, and $GHI$ are each parallel to four parent faces $ABD$, $BCD$, $CAD$, and $CBA$. Thus, in the coarsened tetrahedron, a parent face takes the color of the child face to which it is parallel. This returns the original coloring of the parent tetrahdron before refinement.

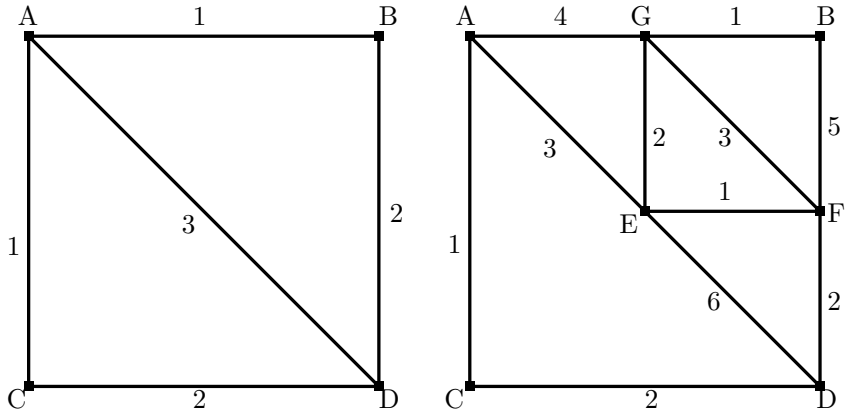


Figure 6: Original (left) and refined (right) meshes with colors.

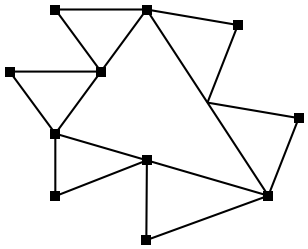

Figure 7: The maximum number of refined neighbors a coarse triangle (center) can have is six.

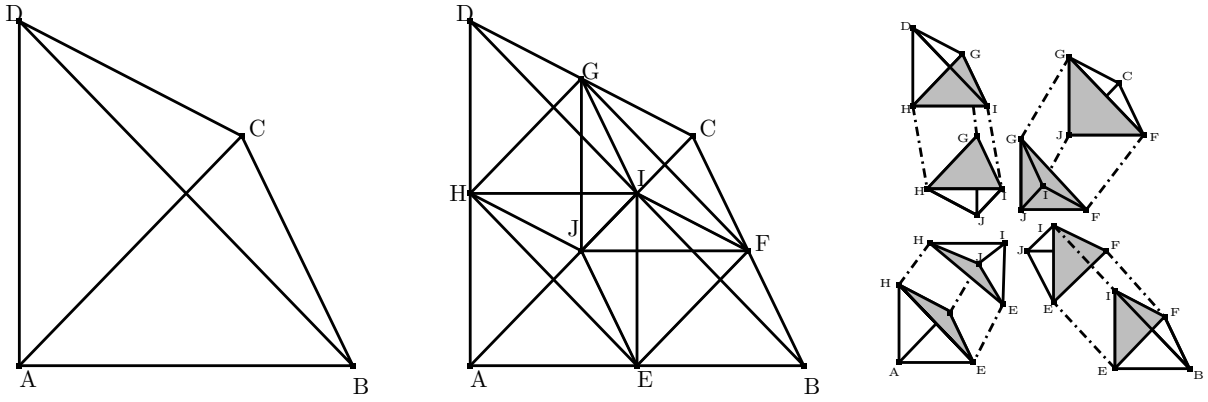


Figure 8: Tetrahedron $ABCD$ (left) is refined into eight new tetrahedra (middle). The image on the right visualizes the small tetrahedra by separating them. The grey faces connected by dashed lines touch.

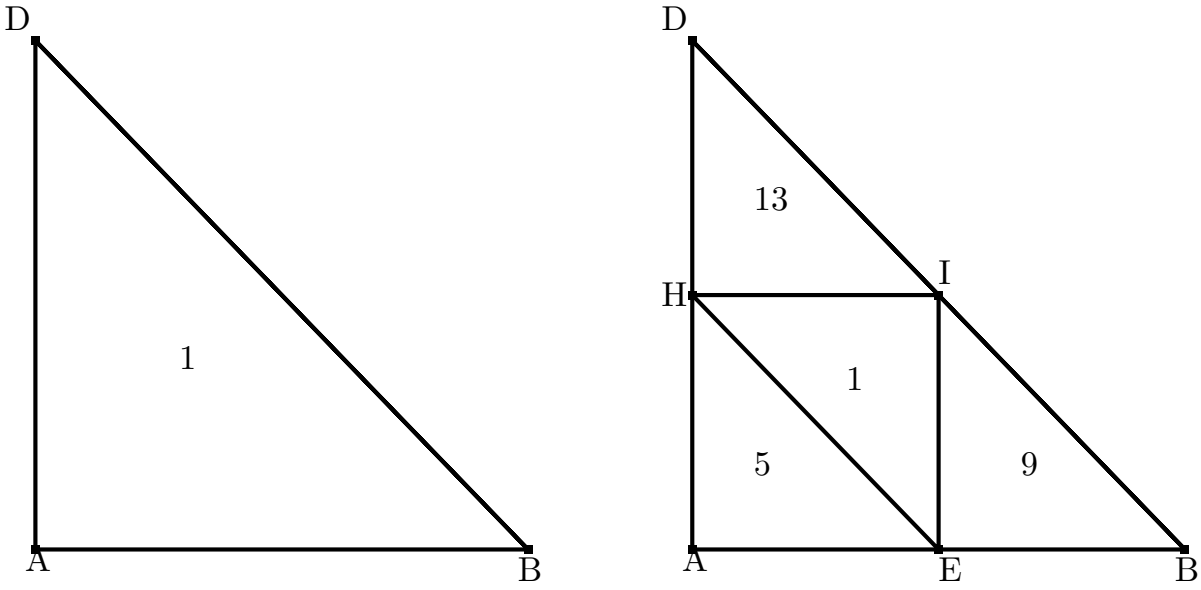


Figure 9: Left: face $ABD$ of tetrahedron $ABCD$ in Figure 8; right: refinement of $ABD$ into four new faces.

# 4 Mesh coloring examples

In Tables 1, 2 and 3, we present a number of examples on meshes of different sizes and geometries for the partitioning algorithm. The coloring algorithm successfully assigns approximately the same number of edges to each available color in a reasonable amount of time. All coloring examples were executed on an Intel Xeon CPU E5530 2.40GHz processor.

For triangular and quadrilateral elements, the meshes are of simple rectangular domains. We consider a sequence of meshes A-E, with A being the coarsest and E being the finest. The meshes are not obtained through nested refinement, though the ratio of the number of elements in consecutive meshes is approximately 4. Additionally, we discretize a spherical shell, i.e. a mesh without boundaries, and a domain containing the NACA-0012 airfoil i.e. a mesh with a cavity. Finally, we color a hybrid quadrilateral/triangular element mesh. Although there are triangles present in the mesh, we use four colors. For tetrahedral elements, meshes A-C are prismatic domains along with a domain containing a three-dimensional extrusion of the NACA-0012 airfoil. We plot some of the meshes considered in Figure 10.

In Figure 11, we plot two metrics that measure the performance of the coloring algorithm. On the left, we display the number of conflicts present in meshes A-E for the different element geometries (triangles, quadrilaterals, and tetrahedra) after the modified greedy coloring. We note that the number of conflicts seems to scale linearly and depend only on the number of surfaces, i.e., edges or faces, in the mesh. On the right of Figure 11, we plot the time it takes these same meshes to be colored (both modified greedy coloring and conflict resolution) according to the two algorithms in Section 3. The coloring runtime appears to scale linearly as well.

| Triangular mesh | No. elements/sides | No. edges in each color | | | Coloring time (s) |
|---|---|---|---|---|---|
| A | 9,752/14,775 | 4,924 | 4,923 | 4,928 | 0.52 |
| B | 39,502/59,547 | 19,845 | 19,845 | 19,857 | 2.58 |
| C | 157,190/236,372 | 78,809 | 78,789 | 78,774 | 12.52 |
| D | 627,326/942,162 | 314,069 | 314,058 | 314,035 | 65.19 |
| E | 2,514,546/3,774,165 | 1,258,087 | 1,258,027 | 1,258,051 | 326.05 |
| NACA-0012 | 2,150/3,278 | 1,095 | 1,092 | 1,091 | 0.08 |
| Sphere A | 3,024/4,536 | 1,512 | 1,512 | 1,512 | 0.02 |
| Sphere B | 189,152/283,728 | 94,576 | 94,576 | 94,576 | 2.06 |

Table 1: Triangular mesh statistics; meshes A-E are on rectangular domains. The NACA-0012 and spherical mesh A are displayed in Figure 10.

| Quadrangular Mesh | No. elements/sides | No. edges in each color | | | | Coloring time (s) |
|---|---|---|---|---|---|---|
| A | 4,373/8,888 | 2,222 | 2,230 | 2,221 | 2,215 | 0.41 |
| B | 17,556/35,396 | 8,845 | 8,852 | 8,851 | 8,848 | 2.48 |
| C | 70,255/141,076 | 35,279 | 35,271 | 35,256 | 35,270 | 11.62 |
| D | 280,988/563,108 | 140,794 | 140,762 | 140,779 | 140,773 | 58.42 |
| E | 1,123,449/2,249,162 | 562,303 | 562,291 | 562,297 | 562,271 | 289.30 |
| Hybrid tri/quad | 3,212/5,132 | 1,273 | 1,283 | 1,285 | 1,291 | 0.05 |

Table 2: Quadrangular and hybrid mesh statistics; meshes A-E are on rectangular domains. The hybrid tri/quad domain is displayed in Figure 10.

| Tetrahedral Mesh | No. elements/sides | No. edges in each color | | | | Coloring time (s) |
|---|---|---|---|---|---|---|
| A | 6,832/14,937 | 3,744 | 3,732 | 3,734 | 3,727 | 0.23 |
| B | 48,879/102,438 | 25,681 | 25,681 | 25,577 | 25,596 | 2.48 |
| C | 351,225/720,736 | 180,211 | 180,234 | 180,232 | 180,059 | 25.70 |
| NACA-0012 | 471,190/973,741 | 243,625 | 243,513 | 243,371 | 243,232 | 30.38 |

Table 3: Tetrahedral mesh statistics; meshes A-C are on prismatic domains. Prismatic domain A is displayed in Figure 10. The NACA-0012 domain is a three-dimensional extrusion of the two-dimensional one in Figure 10.

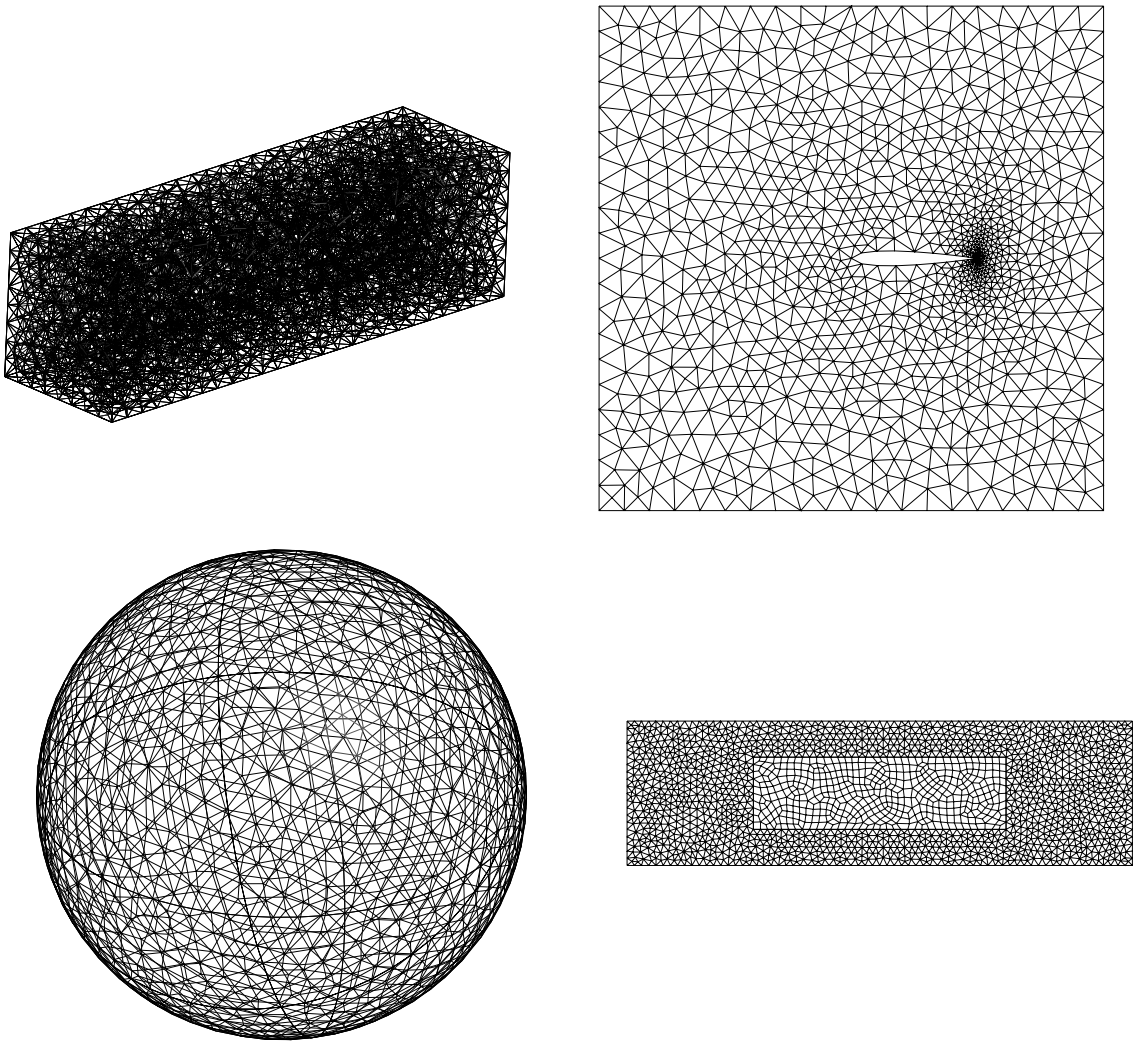

Figure 10: Meshes clockwise from the top left: 3D prism A, two-dimensional NACA-0012 airfoil, rectangular domain of a hybrid hybrid tri/quad mesh, and spherical shell A.

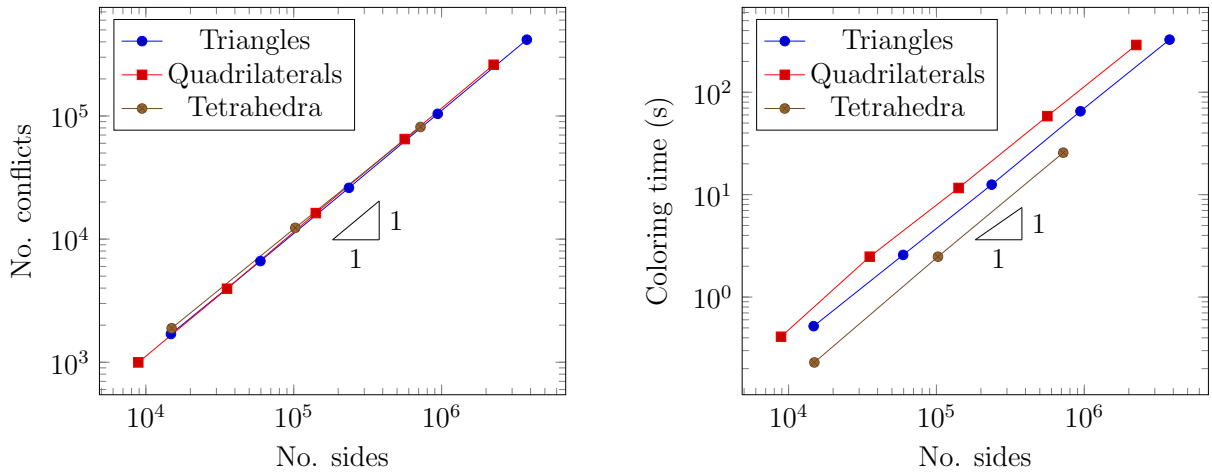


Figure 11: The number of conflicts vs. the number of sides in meshes A-E after the modified greedy coloring for different element geometries (left). Coloring time for the same meshes (right).

| Order of approximation $p$ | 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|---|
| Memory saved (GB) | 0.72 | 1.44 | 2.41 | 3.62 | 5.07 |

Table 4: Memory saved from removing the buffer needed for the mesh E of 2,514,546/3,774,165 (elements/sides) of triangular elements for orders of approximation 1 to 5 for the Euler equations.

### 4.1 Memory savings

By coloring the edges of the mesh, a substantial amount of memory can be saved. If the edges were not colored, two buffers for the surface contribution to the left and right elements would have to be allocated. The necessary memory for these buffers in double precision can be calculated with the formula $2\times(N_p$ basis functions $)\times(N_{\text{eq}}$ equations $)\times(N_{\text{s}}$ surfaces $)\times(8$ bytes$)$. Two buffers are needed because there are two different surface contributions to the left and right elements of a surface in the DG method. For the Euler equations, we show in Table 4 the memory saved on mesh E of triangular elements for various orders of approximation $p$. On GPUs, this is a non-negligible amount of memory, reaching 5 GBs for $p = 5$.

## 5 Data ordering

On GPUs, optimal memory throughput is achieved when sequential threads access sequential locations in memory [17]. Coloring the edges with a minimal number of colors yields a straightforward procedure for ordering the elements and surfaces such that coalesced accesses are favoured. We know that the elements in arbitrary unstructured grids cannot be ordered such that all memory accesses are coalesced.

The parallel DG-GPU flow solver described in [2] assigns one thread per edge in the surface computation kernel. Each thread loads the data for left and right elements

| Edge | Left Element | Right Element |
|---|---|---|
| 0 | 54 | 32 |
| 1 | 8 | 17 |
| ⋮ | ⋮ | ⋮ |
| $N_1 - 2$ | 25 | 12 |
| $N_1 - 1$ | 2 | 13 |

a)

| Edge | Left Element | Right Element |
|---|---|---|
| 0 | 0 | $N_1$ |
| 1 | 1 | $N_1 + 1$ |
| ⋮ | ⋮ | ⋮ |
| $N_1 - 2$ | $N_1 - 2$ | $2N_1 - 2$ |
| $N_1 - 1$ | $N_1 - 1$ | $2N_1 - 1$ |

b)

Table 5: a) arbitrary unstructured edgewise connectivity data for edges in color 1, where $N_1$ is the number of edges of color 1. b) the left and right elements of the edges of color 1 are renumbered such that they are consecutive in memory with respect to the ordering of edges in color 1.

| Edge | Left Element | Right Element |
|---|---|---|
| $N_1$ | 103 | 65 |
| $N_1 + 1$ | 0 | 43 |
| ⋮ | ⋮ | ⋮ |
| $N_2 - 2$ | 34 | 47 |
| $N_2 - 1$ | 2 | 209 |

a)

| Edge | Left Element | Right Element |
|---|---|---|
| $N_1$ | 0 | 43 |
| $N_1 + 1$ | 2 | 209 |
| ⋮ | ⋮ | ⋮ |
| $N_2 - 2$ | 102 | 31 |
| $N_2 - 1$ | 103 | 65 |

b)

Table 6: a) connectivity data for edges in color 2, where $N_2$ is the number of edges of color 2. b) edges are reordered with ascending left element.

that share the edge. These memory accesses for edges will be coalesced if sequential edges have sequential left and right elements, i.e., the $k$th thread loads data for the left element at memory location `offset1 + k` and data for the right element at memory location `offset2 + k`. Without element and edge reordering, these accesses will likely be uncoalesced due to the irregular nature of unstructured meshes.

Mesh connectivity is stored as pointers: an edge points to the two elements that share it. Table 5 a) shows an excerpt of edgewise connectivity data required for an unstructured mesh. The $k$th row stores data for edge $e_k$. The first and second columns store the element numbers of each edge's left and right element. For example, $e_0$ has left element $\Omega_{54}$ and right element $\Omega_{32}$.

We propose the following element and edge renumbering method to reduce occurrences of irregular access patterns. This is done as a mesh preprocessing stage. First, color the edges of the mesh and reorder them based on their assigned color, i.e. edges 0 to $N_1$ are of color 1, where $N_1$ is the number of edges of color 1, followed by edges of color 2 and 3. That is, we have a connectivity table as shown in Table 5 a) for color 1.

Next, the left and right elements of the $k$th edge in color 1 are renumbered as element $k$, and $N_1 + k$, respectively, where $N_1$ is the number of edges in color 1. Table 5 b) represents the renumbered connectivity data for the edges in color 1. Because every element has exactly one edge of color 1, all the elements in the mesh are renumbered.

This procedure yields entirely coalesced acccesses for the edges in color 1, i.e. for

| 5 colors | 3 colors | Element renumbering | Edge reordering | Solver runtime (s) |
|---|---|---|---|---|
| x | | | | 166.79 |
| | x | | | 158.14 (5.1) |
| x | | x | | 177.93 |
| | x | x | | 153.31 (8.0) |
| x | | x | x | 144.75 (13.2) |
| | x | x | x | 129.34 (22.4) |

Table 7: Execution runtimes for 1000 timesteps of the double Mach test problem with $p = 1$ on mesh D presented in Table 1 for different data reorganization strategies. The number in parentheses is the percent speed-up relative to the unordered, unrenumbered mesh of five colors.

| Mesh | A | B | C | D | E |
|---|---|---|---|---|---|
| Time (s) | 0.08 | 0.37 | 1.81 | 7.98 | 33.26 |

Table 8: Total renumbering and reordering times for the two-dimensional meshes A-E of triangles with three colors.

approximately a third of the edges for triangles and a fourth of the edges for quadrangles and tetrahedra.

Because all the elements have been renumbered, the order of the left and right elements of edges in the remaining color groups may not be consecutive, as shown in Table 6 a). We can no longer renumber elements, but we can reorder edges in these groups so that the left element numbers are in ascending order, see Table 6 b). Therefore some memory transfers for left elements will be coalesced.

Reorganizing the mesh can be done efficiently using hash tables even for problems of substantial size. For example, the total element reordering and edge renumbering time of mesh D with three colors presented in Table 1 (627,326 elements, and 942,162 sides) is 7.98 seconds. For the other meshes of triangular elements (meshes A-E), see Table 8.

## 5.1 Data ordering examples

We now apply this element and edge renumbering scheme to the DG-GPU solver for the double Mach reflection test problem described in [2] on an NVIDIA GTX 580 graphics processing unit, which has 3 GB of video memory. We run the solver with a polynomial approximation $p = 1$ for 1000 timesteps on mesh D with triangular elements referenced in Table 1. In Table 7, we measure the solver performance by comparing the runtimes of a mesh colored with an optimal three colors and meshes with a suboptimal five colors. Compared with a mesh of five colors without element renumbering or edge reordering, we find that reducing the number of colors gives minor impovement ($\sim$ 5 % speed-up). These savings result from less overhead from kernel launches [7]. We also report the runtime for element-renumbered meshes that were colored with three and five colors. Here, using more colors than required actually increases runtime. Next, we combine

element renumbering and edge reordering. With five colors, we have an improvement of $\sim$ 13 %. The speed-up increases to approximately 22 % when the renumbered and reordered mesh has three colors.

# 6 Conclusion

We have presented a simple method to mitigate the race condition that arises during the surface integral evaluation of shared-memory-based parallel DG and finite volume type flow solvers. Although we do not test on a finite volume solver, our approach is applicable to this method as well. The main motivation of this work was to reduce the memory footprint of the code by eliminating buffers [2, 3, 4]. This procedure colors the edges of unstructured triangular, tetrahedral and quadrilateral meshes using an optimal number of colors; as a result, the number of kernel launches is reduced, contributing to code simplicity. We do not have a proof that the coloring algorithm always terminates. However, we have not encountered a mesh for which it failed to do so. We note that practical edge coloring algorithms existing in literature [14, 15] do not have formal proofs either.

We have proposed a coloring algorithm extension to adaptive mesh refinement. The original coloring can be done in a preprocessing stage and adaptive coloring is a simple mapping from the original coloring, which can be done quickly and in parallel.

The same coloring is used for reorganizing the DOFs in memory to reduce access latencies. The reordering scheme allows for approximately a 22 % speed-up when applied to a DG solver of the Euler equations in two-dimensions.

The edge coloring algorithm can also serve as an approach to quadrilateral mesh generation. Deleting all the edges of one chosen color creates a mesh comprised of quadrilaterals everywhere in the domain interior. This is similar to Blossom-Quad [18], a perfect matching algorithm. It will be interesting to investigate if our coloring algorithm yields meshes of comparable quality in a similar amount of time. We leave this for a future work.

# 7 Acknowledgements

This work was supported in part by the Natural Sciences and Engineering Research Council of Canada grant 341373-07, and Alexander Graham Bell PGS-D grant.

# References

[1] M. Harris, "Using shared memory in CUDA C/C++." Accessed online: http://devblogs.nvidia.com/parallelforall/using-shared-memory-cuda-cc/ on August 29, 2015.

[2] M. Fuhry, A. Giuliani, and L. Krivodonova, "Discontinuous Galerkin methods on graphics processing units for nonlinear hyperbolic conservation laws," *International Journal for Numerical Methods in Fluids*, vol. 76, no. 12, pp. 982–1003, 2014.

[3] A. Klöckner, T. Warburton, J. Bridge, and J. Hesthaven, "Nodal discontinuous Galerkin methods on graphics processors," *Journal of Computational Physics*, vol. 228, no. 21, pp. 7863 – 7882, 2009.

[4] A. Corrigan, F. F. Camelli, R. Löhner, and J. Wallin, "Running unstructured grid-based CFD solvers on modern graphics hardware," *International Journal for Numerical Methods in Fluids*, vol. 66, no. 2, pp. 221–229, 2011.

[5] Y. Xia, L. Luo, J. Lou, H. Luo, and J. Edwards, "Development of a portable, GPU-accelerated high-order discontinuous Galerkin CFD code for compressible flows on hybrid grids." Accessed online: http://synergy.cs.vt.edu/afosr-bri/files/workshop-2014/02-RDGFLO3D.pdf on August 29, 2015.

[6] R. Löhner, *Applied Computational Fluid Dynamics Techniques: An Introduction Based on Finite Element Methods*. Wiley, 2008.

[7] Y. Xia, J. Lou, H. Luo, J. Edwards, and F. Mueller, "OpenACC acceleration of an unstructured CFD solver based on a reconstructed discontinuous Galerkin method for compressible flows," *International Journal for Numerical Methods in Fluids*, vol. 78, no. 3, pp. 123–139, 2015.

[8] D. Komatitsch, D. Michéa, and G. Erlebacher, "Porting a high-order finite-element earthquake modeling application to NVIDIA graphics cards using CUDA," *J. Parallel Distrib. Comput.*, vol. 69, pp. 451–460, May 2009.

[9] M. J. Aftosmis, M. J. Berger, and S. M. Murman, "Applications of space-filling curves to cartesian methods for cfd," *AIAA Paper*, vol. 1232, p. 2004, 2004.

[10] D. Burgess and M. Giles, "Renumbering unstructured grids to improve the performance of codes on hierarchical memory machines," *Advances in Engineering Software*, vol. 28, no. 3, pp. 189–201, 1997.

[11] V. G. Vizing, "On an estimate of the chromatic class of a $p$-graph," *Diskret. Analiz No.*, vol. 3, pp. 25–30, 1964.

[12] R. Wilson, *Introduction to Graph Theory*. Longman, 1996.

[13] K. Appel and W. Haken, "Every planar map is four colorable. I. Discharging," *Illinois J. Math.*, vol. 21, no. 3, pp. 429–490, 1977.

[14] T. T. Lee, Y. Wan, and H. Guan, "Randomized $\delta$-edge colouring via exchanges of complex colours," *International Journal of Computer Mathematics*, vol. 90, no. 2, pp. 228–245, 2013.

[15] M. Fiol and J. Vilaltella, "A simple and fast heuristic algorithm for edge-coloring of graphs," *arXiv preprint arXiv:1210.5176*, 2012.

[16] W. Bangerth, R. Hartmann, and G. Kanschat, "deal.II – a general purpose object oriented finite element library," *ACM Trans. Math. Softw.*, vol. 33, no. 4, pp. 24/1–24/27, 2007.

[17] NVIDIA, *CUDA C Programming Guide v7.5*. Last updated September 1, 2015.

[18] J. Remacle, J. Lambrechts, B. Seny, E. Marchandise, A. Johnen, and C. Geuzainet, "Blossom-quad: A non-uniform quadrilateral mesh generator using a minimum-cost perfect-matching algorithm," *International Journal for Numerical Methods in Engineering*, vol. 89, no. 9, pp. 1102–1119, 2012.